

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 06-07-2005		2. REPORT TYPE Final Report		3. DATES COVERED (From – To) 21 June 2004 - 21-Jun-05	
4. TITLE AND SUBTITLE Catalytic Reaction Sets: A First-Principle Approach To The Evolution Of Communication			5a. CONTRACT NUMBER FA8655-04-1-3045		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Professor Jose F Fontanari			5d. PROJECT NUMBER		
			5d. TASK NUMBER		
			5e. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Universidade de Sao Paulo Av. Trabalhador Sao Carlense 400 Caixa Postal 369 Sao Carlos CEP 13560-970 Brazil				8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) EOARD PSC 802 BOX 14 FPO 09499-0014				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) Grant 04-3045	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This report results from a contract tasking Universidade de Sao Paulo as follows: The Grantee will investigate how the information processing capability of an agent (i.e., the learning algorithm the agent uses to infer the meaning of a signal it receives and to decide which signal to emit given the environment variables) determines the structure of the language that emerges in the population.					
15. SUBJECT TERMS EOARD, Communications, Command and Control, Information Integrity, Information Assurance					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UL	18, NUMBER OF PAGES 38	19a. NAME OF RESPONSIBLE PERSON VALERIE E. MARTINDALE, Lt Col, USAF
a. REPORT UNCLAS	b. ABSTRACT UNCLAS	c. THIS PAGE UNCLAS			19b. TELEPHONE NUMBER (Include area code) +44 (0)20 7514 4437

Final Report

**Catalytic reaction sets: a first-principle approach to the
evolution of communication**

Grant FA8655-04-1-3045

José Fernando Fontanari
Instituto de Física de São Carlos, Universidade de São Paulo,
Caixa Postal 369
13560-970 São Carlos SP Brazil
Tel: +55-16-33739849, Fax: +55-16-33739877, e-mail: fontanari@ifsc.usp.br

Contents

Research activities	3
The “true” number of objects in the world: Akaike Information Criterion	3
Categorization of complex objects	8
Discussion	17
Compositional communication codes in the synthetic ethology framework	18
Use of the award resources	20
Acknowledgement of Sponsorship	21
Disclaimer	21
Disclosure of inventions	21

1 Research activities

The main results of the research activities supported by EOARD were described in great detail and made public in the three papers published in the Proceedings of the 2005 International Conference on Integration of Knowledge Intensive Multi-Agent Systems (KIMAS 2005), Eds. C. Thompson and H. Hexmoor, ISBN 0-7803-9013-X:

1. “Evolution of communication in a community of simple-minded agents”, presented in session WM4: Evolution of Communication and Cognition, pp. 285-290 of the Proceedings of KIMAS’05.
2. “Minimal Models for Text Production and Zipf’s Law”, presented in session WM4: Evolution of Communication and Cognition, pp. 297-300 of the Proceedings of KIMAS’05.
3. “Meaning Creation and Modeling Field Theory”, presented in session WA3: Semiotics, Language and Meanings, pp. 405-410 of the Proceedings of KIMAS’05.

These papers addressed the main topics of investigation listed in the original proposal and for the sake of completeness they are attached to this report (see paper_1.pdf, paper_2.pdf, and paper_3.pdf).

The results presented in the sequel are natural extensions of the research described in the Interim Report and will soon be submitted to publication.

1.1 The “true” number of objects in the world: Akaike Information Criterion

To instantiate any model of communication between virtual or real organisms, a basic cognitive requirement must be fulfilled, namely, that the organisms be capable of classifying different types of situations and, accordingly, be capable of recognizing that a situation of a particular type turns up. The effectiveness of the Modeling Field Theory, MFT for short, framework as an autonomous mechanism for the spontaneous formation of meaning or, equivalently, for category creation has already been demonstrated in the previous report. Here we use the same simple one-dimensional environment, originally proposed by Luc Steel¹, in which an organism inhabits an abstract world made up of N objects or situations, each of which described by a single feature value modeled by a real variable $O_i \in (0,1), i = 1, \dots, N$ drawn from some probability distribution. These features are, of course, abstract and have no particular meaning in the model, though it may be helpful to think of them as perceptual features such as color or smell. The question is whether such organism is capable to produce a repertoire of features to succeed in discriminating among

¹ L. Steels, “Perceptually grounded meaning creation,” Proceedings of the Second International Conference on Multi-Agent Systems, ICMAS-96, 338-344, 1996.

the known objects and to adapt that repertoire when new objects are incorporated into the environment.

In the MFT scheme we assume that there are M concept-models described by real-valued variables $S_k, k = 1, \dots, M$ that should represent the objects $O_i, i = 1, \dots, N$. We define arbitrarily the following partial similarity measure between object i and concept k

$$l(i | k) = (2\pi\sigma_k)^{-1/2} \exp\left[-(O_i - S_k)^2 / 2\sigma_k^2\right], \quad (1)$$

where, at this stage, the fuzziness σ_k is a parameter given *a priori*. The goal is to find an assignment between models and objects such that the global similarity

$$L = \frac{1}{M} \sum_i \log \sum_k l(i | k) \quad (2)$$

is maximized. For our purposes, namely, to compare the values of L obtained using distinct number of fields, it is *very important* that we re-normalize the global similarity by the number of fields, as done in Eq. (2), in order to make it an intensive quantity with respect to M . This maximization can be achieved using the MFT mechanism of concept formation which is based on the following dynamics for the modeling fields²

$$dS_k/dt = \sum_i f(k | i) [\partial \log l(i | k) / \partial S_k] \quad (3)$$

where the fuzzy association variables $f(k | i)$ are defined by

$$f(k | i) = l(i | k) / \sum_{k'} l(i | k') \quad (4)$$

and give a measure of the correspondence between object i and concept k relative to all other concepts k' . In fact, it can be shown that this dynamics always converges to a (usually local) maximum of the similarity L . However, by properly adjusting the fuzziness σ_k the global maximum can be singled out. In particular, here we choose to decrease the fuzziness on the flight, i.e., during the time evolution of the modeling fields according to the following prescription

$$\sigma_k^2(t) = \sigma_{k1}^2 \exp(-\alpha t) + \sigma_{k0}^2 \quad (5)$$

with $\alpha = 5 \times 10^{-4}$, $\sigma_{k1} = 1 \forall k$ and $\sigma_{k0} = 0.03 \forall k$. We have shown that this setting allows perfect categorization, in a sense that the values of the modeling fields match those of the objects, provided that the number of modeling fields M is equal or greater than the number of objects N .

This framework, however, does not account for a need to decide how many different models (i.e., modeling fields) the organism really needs. A biological organism evolves various complex mechanisms, related to instinctual and emotional evaluations, to make such a decision, i.e., to distinguish between the objects and the meaningless background that compose its world. An adaptation of a quote by Ferdinand de Saussure may be

² L. I. Perlovsky, *Neural Networks and Intellect: Using Model-Based Concepts*, Oxford: Oxford University Press, 2001.

appropriate to describe the situation – without labels the world is a vague, uncharted nebula³. But too many labels is equivalent to have no labels at all. In fact, mathematical approaches to determine the true number of objects are nontrivial because any data can be better fitted with more models (i.e., concepts). Therefore it is necessary to balance maximization of similarity, Eq. (2), against the number of parameters in the model. A theoretically consistent way to achieve this balance is to use Akaike Information Criterion, AIC for short, which is an asymptotic correction to the similarity function related to the bias due to the number of parameters⁴, namely,

$$AIC = L - \frac{1}{2} M_{par} \quad (6)$$

where M_{par} is the number of adjustable parameters of the models and L is given by Eq. (2).

In our case, since there are two parameters per model (S_k and σ_k) we have $M_{par} = 2M$.

To better appreciate the effectiveness of the AIC to single out the true number of objects in the environment we consider a very simple situation in which there are $N = 4$ objects: $O_1 = 0.2$, $O_2 = 0.4$, $O_3 = 0.6$ and $O_4 = 0.8$. The modeling field dynamic equations (3) – (5) are then solved numerically with Euler's method using the step-size $h = 10^{-4}$ for several choices of M and the resulting value of the AIC, as given by Eq. (6), is plotted against time t . The results shown in figure 1 illustrate how tricky the determination of the true value of N can be. Indeed, for short times, the choice of fewer models than the true number yields the maximum value of AIC, but as the dynamics progresses the insufficiency of models becomes readily noticeable and, as expected, in the asymptotic regime $t \rightarrow \infty$ the maximum of AIC corresponds to the situation $M = N$. Interestingly, the observed decrease of AIC in the unrealizable case $M < N$ yields a clear indication that something is going wrong, serving thus as a warning to increase the number of models. On the other hand, by following the evolution in over-realizable case $M > N$, say $M = 6$, we find no signs that we are using superfluous models.

Taking advantage of the distinctive behavior pattern of the dependence of AIC on t in the unrealizable case, we envisage a simple strategy to adjust the value of M on the flight: starting with a single model S_1 , we create a new model whenever AIC decreases. The value of the new modeling field created at $t = t_c$, say $S_2(t_c)$, is then given by a perturbation of one of the previous fields, e.g., $S_2(t_c) = S_1(t_c) + 0.01\varepsilon$, where ε is a random number drawn uniformly in the interval $(-1,1)$. In addition, the fuzziness of the new model obeys the rescaled equation (5), $\sigma_2^2(t) = \sigma_{21}^2 \exp[-\alpha(t - t_c)] + \sigma_{20}^2$. The trouble with this procedure is that by adding a new model that, in principle, has a small similarity with all objects, we simultaneously decrease L and increase M_{par} in Eq. (6), which results in a further decrease of AIC. To circumvent this difficulty we must allow some time, i.e., a time interval

³ The original quote is “Without language, thought is a vague, uncharted nebula. There are no pre-existing ideas, and nothing is distinct before the appearance of language”, in de Saussure, F. 1966. Course in General Linguistics. Translated by Wade Baskin. New York: McGraw-Hill Book Company.

⁴ H. Akaike, Statistical predictor identification. Ann. Stat. Math. **22**, 203-217, 1970.

$\Delta t = 3000$, for the new field to adapt to the objects and only then to check for a decrease of AIC. The result of applying this strategy to the same categorization




Figure 1 Illustration of the use of Akaike Information Criterion (AIC) measure in conjunction with the MFT scheme with $M = 2, 3, 4, 5$ and 6 modeling fields to determine the number of objects in the environment. Here the true number is $N = 4$, which corresponds to the maximum of the AIC for large t .

problem addressed in figure 1 is depicted in figure 2: it is clearly a success! Details of the time evolution of the modeling fields are presented in figure 3 (we arbitrarily assign the value 0 to the dormant modeling fields).




Figure 2 Results of the adaptive scheme to find the true number of objects for the same problem of figure 1. Starting with a single model ($M=1$) the evolution of AIC measure is followed until a decrease is detected (this check is done at time intervals of $\Delta t = 3000$) then a new model is created. The arrows indicate the moments when a new model is created.




Figure 3 Time evolution of the modeling fields using the adaptive scheme to create new fields on the flight based on the behavior pattern of the AIC. These data correspond to the same experiment depicted in the previous figure.

1.2 Categorization of complex objects

Up to now we have considered the objects as points on a single axis. In this section we assume that an object is a set of points drawn from a Gaussian distribution with mean m and variance v^2 . The issue here is to verify what conditions need to be satisfied in order that the MFT system recognizes the whole object and not the individual points that compose it. Of course, we expect that the final categorization ability of the system will depend strongly on the balance between the baseline resolution of the modeling fields σ_{k0}^2 , the variance v^2 and the distance between the means of the distributions associated to each object. In figures 4, 5 and 6 we illustrate the performance of the MFT scheme to categorize complex objects that do not overlap. In this sense this problem is not much different from that of the “simple” objects (i.e., single points) discussed before, and so one might think that this may be the reason for the good performance of our approach in this case as well. The case in which two objects overlap are considered in figures 7 and 8, where we assume for simplicity that all objects standard deviations v are the same and equal to the baseline standard deviation of the models σ_{k0} . A more challenging case with four overlapping objects is presented in figures 9 and 10. Details and discussion are presented in the figure captions.




Figure 4 Akaike Information criterion measure in the case that 20 points were generated following a Gaussian distribution of mean $m = 0.5$ and standard deviation $\nu = 0.01$. The baseline standard deviation of the modeling fields is $\sigma_{k0} = 0.03$. Maximization of AIC yields the correct answer, namely, there is only one (complex) object in the environment.




Figure 5 Akaike Information criterion measure in the case 40 points were generated following two Gaussian distributions (20 points for each object) of mean $m_1 = 0.3$ and $m_2 = 0.6$, and standard deviations $v_1 = v_2 = 0.01$. As before, the baseline standard deviation of the modeling fields is $\sigma_{k0} = 0.03$. Maximization of AIC yields the correct answer, namely, there are two objects in the environment. Note the pronounced decrease of the AIC measure for $M = 1$ at large t , similarly to our findings in the case the objects were single points.




Figure 6 Modeling fields for the case $M = 2$ of the previous figure. The points representing the two objects are shown in blue. In particular $\nu = 0.01$ for all objects and $\sigma_{k0} = 0.03$ for all models.




Figure 7 Akaike Information criterion measure in the case 100 points were generated following two Gaussian distributions (50 points for each object) of mean $m_1 = 0.3$ and $m_2 = 0.6$, and standard deviations $v_1 = v_2 = 0.2$. The baseline standard deviation of the modeling fields is $\sigma_{k0} = 0.2$ too. Now, maximization of AIC does not yield the correct answer, but considering the difficulty of the problem (see next figure) the prediction of $M = 3$ followed closely by the correct solution $M = 2$ is not bad at all. The AIC measure for $M = 1$ is not shown because it is too small and does not fit in the scale of the figure.




Figure 8 Modeling fields for the case $M = 2$ of the previous figure. The points representing the two objects are shown in blue. Although the environment consists of two Gaussian objects centered at 0.3 and 0.6, this solution does not correspond to the maximum of the AIC measure. However, a similar plot of the modeling fields for the maximum $M = 3$ indicates that the system still uses only two fields (i.e., $S_1 = S_2 \neq S_3$) but with slightly different values.




Figure 9 A tough problem: the environment is composed of four objects each of which represented by 100 points drawn from Gaussian distributions of means 0.2, 0.4, 0.6, and 0.8, and standard deviation $\sigma = 0.2$. The 400 points are plotted in the figure, one symbol for each object. The symbols are shown displaced vertically, four symbols per row, for ease of visualization. The original data is recovered by projecting all symbols in a single row. Would the reader be able to tell how many objects there are in the figure, if they were plotted with the same symbol? This is the task we set to our MFT system.




Figure 10 Akaike Information criterion measure for the problem stated in figure 9. The baseline standard deviation of the modeling fields is $\sigma_{k0} = 0.08$. Surprisingly, maximization of the AIC measure *for large t* yields the correct answer $M = 4$. However, the time dependence of this measure is very different from that observed in the simpler problems analyzed in figures 1, 4, 5 and 7. In particular, there is a transient stage when the AIC measure increases until it reaches a maximum and then decreases towards a fixed value. This odd behavior pattern, due to a “problem” in our theoretical formulation which is discussed in section 1.3, jeopardizes completely the automated scheme for generating new models we used to draw figures 2 and 3.




Figure 11 Modeling fields for the case $M = 4$ of the problem stated in figure 9. The points representing the four objects are shown in blue. The data correspond to the AIC measure of the previous figure. It is amazing that the system can actually single out four objects in the cloud of points depicted in figure 9.

1.3 Discussion

Looking at the time dependence of the AIC measure for *fixed* M , depicted in figures 1, 4, 5, 7 and 10, immediately brings a question up: Shouldn't L (or, equivalently, AIC if M is kept fixed) be a increasing function of time? Yes, provided the fuzziness $\sigma_k, k = 1, \dots, M$ is kept fixed during the evolution of the fields S_k , which is not the procedure we are adopting here since Eq. (5) provides an explicit prescription for updating the fuzziness. Hence there is actually no reason to expect that L or the AIC measure will increase during the time evolution of the fields. This may be a problem if we try to find the true number of objects by maximizing the AIC measure with respect to M since the optimum value of M may depend on the instant of time we look at the fields (see figure 10). Of course, the ambiguity is resolved if one accords that the maximization is carried out in the asymptotic limit (large t) only, but this amounts to discard a solution that has a higher value of the AIC measure (e.g., the set of fields at $t/50 = 25$ in figure 11). Is this satisfactory? To answer this question, let us ask another one: is there a way to update the fuzziness so as to guarantee that L increases with increasing t ? Yes, considering σ_k as an adjustable parameter, similar to the modeling fields, we obtain the equation

$$d\sigma_k/dt = \sum_i f(k|i) [\partial \log l(i|k) / \partial \sigma_k] \quad (7)$$

which solved simultaneously with Eq. (3) leads to the maximum of L . We have solved this set of $2M$ coupled equations and the result was almost always the uniform solution, i.e., $S_1 = S_2 = \dots = S_M$ and $\sigma_1 = \sigma_2 = \dots = \sigma_M$. In fact, looking again at figure 11 we can see that considering a single field with a large standard deviation can account for most of the points in the environment – this is the optimal, but unsatisfactory, solution for a difficult problem such as that posed in figure 9. In our setting the homogenous solution breaks down (7) would allow the fuzziness to remain at a large value] so the single field can no longer account for all points in the environment, resulting in the decrease of the AIC measure.

In summary, although the combination of the MFT scheme and the AIC measure does indeed solve some difficult categorization problem (by solving we mean to find the true number of objects and create a suitable representation for them, see figure 11) we have not yet succeeded to produce a neat theoretical framework to describe the combination of those two tools. In particular, the relation between the dynamics given by Eqs. (3) - (5) and the global similarity L is obscure. Perhaps a more consistent approach, to be pursued in the future, is to consider two time-scales for the fields S_k and the fuzziness σ_k : if the latter evolves much slower than the former then it would be correct to say that L is a Lyapunov function of the dynamics. The decrease of σ_k could then be viewed as a procedure similar to the cooling schedule of the Simulated Annealing algorithm. It remains to be seen whether this theoretically more satisfactory framework actually works in practice.

1.4 Compositional communication codes in the synthetic ethology framework

Human language is one of the few biological phenomena that still resist a purely evolutionary explanation as offered by Darwin's concept of evolution through natural selection. In fact, non-human animals communication codes (proto-languages) are typically non-syntactic, i.e., signals refer to whole situations, in contrast to human language which is characterized by signals formed by discrete components that have their own meaning. That composition allows us to take advantage of combinatorics and so as linguists put that to "make infinite use of finite means".

The emergence of compositional syntax has been extensively studied within a framework for modeling the cultural evolution of language - the so-called Iterated Learning Model⁵. There language is seen as a mapping between meaning and signals. Signals are defined as strings of symbols drawn from some alphabet Σ . Meanings are vectors of F components, each of which taking on V discrete values. For example, consider the following "language", i.e. mapping meaning \rightarrow signal, in which the signal strings are of fixed length $l = 3$, and $F = 3$, $V = 2$:

$\{1,2,2\} \rightarrow adf$; $\{1,1,1\} \rightarrow ace$; $\{2,2,2\} \rightarrow bdf$; $\{2,1,1\} \rightarrow bce$; $\{1,2,1\} \rightarrow ade$;
 $\{1,1,2\} \rightarrow acf$.

This language is compositional because a sub-signal (i.e., a part of the signal string) represents a feature value of the meaning vector. In particular, whenever the first entry takes on the value 1 the corresponding signal begins with symbol a , if it takes on the value 2, the signal begins with b . The mapping meaning-signal possesses a structure which can be inferred by the learner to create a unique signal for new meanings such as $\{2,1,2\}$ ⁶. This contrasts with a holistic language for which a random signal is assigned to each meaning. Of course, the proposed meaning-signal mapping can account for both extremes (holistic and compositional) as well as for intermediate languages.

The main problem with the Iterated Learning Model (ILM) is the mind reading assumption: when an agent observes a signal, the intended meaning of that signal is also given. That actually makes communication superfluous and so it is unwise, to say the least, to base the study of cultural evolution in a framework that relies in such odd ability. The computational approach to the evolution of communication based on MacLennan's synthetic ethology⁷ circumvents this difficulty. In fact, the central idea of MacLennan's

⁵ See, e.g., H. Brighton, Compositional Syntax from Cultural Transmission, *Artificial Life* **8**, 25-54, 2002.

⁶ According to the rules of our model language the signal must be bcf .

⁷ B. J. MacLennan, Synthetic ethology: an approach to the study of communication *Artificial Life II*, SFI Studies in the Sciences of Complexity, vol. X, 631-658, Addison-Wesley, 1991.

framework⁸ is that an agent be capable to guess or infer the meaning of a symbol written in a public board by some other agent. There is no place for mind-reading in this scenario.

In what follows we will present preliminary results of our attempt to combine the very interesting meaning-signal mapping described above, which has been extensively employed in the ILM framework, with the synthetic ethology framework described in Fontanari & Perlovsky⁸.

There are N agents that interact among themselves by perceiving and making changes in the environment they inhabit. The environment of each agent is composed of two parts – a public environment shared with all other agents in which the signals are written, and a private environment – the agent mind – to which no other agent has access. The architecture of the agents' world is illustrated in figure 12.




Figure 12 The structure of the world inhabited by four agents. The public environment is used as a blackboard to read, erase and write signals; the private environment is the agent's mind where the meanings are hidden from the other agents.

The public environment can be found a finite number of states, each state represented by the integer $\gamma \in \{1, 2, \dots, G\}$. Actually each value of γ corresponds to a string of L symbols, say, $\omega_1 \omega_2 \dots \omega_L$ so that G depends on the size of the alphabet $|\Sigma|$ and the length L , i.e., $G = |\Sigma|^L$. Similarly, the state of the agents minds is described by the integer variable $\lambda \in \{1, 2, \dots, H\}$, where as before each value of λ labels a vector (m_1, m_2, \dots, m_F) with $m_i = 0, 1, \dots, V-1$. Hence $H = V^F$. The basic idea is to permit the agents to exchange information in the content of their minds by reading from and writing on (i.e., modifying the state) the public environment. Of course, the agents must be endowed with the

⁸ See also J.F Fontanari & L. Perlovsky, Evolution of communication in a community of simple-minded agents, Proceedings of KIMAS'05, 285-290, 2005.

capability to respond actively to the stimuli coming from both internal and external environments. That response is of two types, an *action* and an *emission*. When performing an emission, the agent draws a signal γ' from the set of signals used to describe the public environment and replaces the current state of that environment by γ' . In doing so, the agent modifies the public environment and so emitting is like signaling. On the other hand, an action is a more introspective business: prompted by the signal, say γ , placed in the public environment the agent draws a situation λ' from the private environment repertoire. In other words, the agent interprets the signal γ as meaning situation λ' . Of course, the correctness of this inference will depend on whether the private environment of the agent that last modified the public environment (by writing the symbol γ on it) was λ' . In that case, we say that a successful communication event has occurred and both agents involved in that event are rewarded in the sense of having their fitness increased by one unity of fitness. . In summary, the agents are modeled by finite state machines with two transition tables for each agent

$$A(\gamma, \lambda) \mapsto \lambda' \quad \text{and} \quad E(\gamma, \lambda) \mapsto \gamma' \quad (8)$$

depending on whether the agent is acting or emitting, respectively.

We already know from our previous work⁸ that the genetic algorithm is sufficiently powerful to produce an optimal communication code given the constraints on the size of the signal and meaning spaces. Such a code will most probably be a holistic code, since the entries of the transition tables (8) which ultimately determine the behavior of the agents is not affected by the nontrivial structure of the meaning-signal mapping. Hence our previous formulation of synthetic ethology scenario must be modified in order to allow for the emergence of compositionality. The way to do that is to reward the fitness of the agents involved in “almost” successful communication events, i.e., by assuming the fitness is some decreasing function of the distance between the real meaning λ and the inferred meaning λ' . This way we naturally introduce the metric of the meaning space into the problem. In addition, assuming that mutation takes place on signal only, by modifying not the entire string but only one of its component symbols allows selection to act on the structure of the signal space too. The computational implementation of this scheme is on the way and we plan to submit the result of that study to the Evolution of Language Conference .

2 Use of the award resources

The two \$ 2,700.00 checks corresponding to 2/3 of the funds granted by EOARD were cashed by the Instituto de Física de São Carlos (IFSC) and so incorporated into the Institute budget. The funds were used to buy a notebook (\$ 2,500.00), a Compaq Fortran Compiler (\$ 1,200.00, including custom and mail taxes), a Microsoft Office package (\$ 300.00), toner and ink for printers (\$ 700.00), and books on the evolution of language (\$

400.00). The expenses were carried out according to the strict financial rules of the Universidade de São Paulo (USP) and all goods acquired are now property of the university.

3 Acknowledgement of Sponsorship

Effort sponsored by the Air Force Office of Scientific Research, Air Force Material Command, USAF, under grant number FA8655-04-1-3045. The U.S. Government is authorized to reproduce and distribute reprints for Government purpose notwithstanding any copyright notation thereon.

4 Disclaimer

The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research or the U.S. Government.

5 Disclosure of inventions

I certify that there were no subject inventions to declare during the performance of this grant.

Evolution of communication in a community of simple-minded agents

José F. Fontanari

Universidade de São Paulo, São Carlos, Brazil, fontanari@ifsc.usp.br

Leonid I. Perlovsky

Air Force Research Lab, Hanscom AFB, USA, Leonid.Perlovsky@hanscom.af.mil

Abstract — *We re-examine the seminal work of MacLennan on the evolution of communication in a population of simple agents (finite-state machines). The original model is modified by separating the signaling and the responding systems in two independent modules, which facilitates greatly the analysis of the behavior of the agents. We have carried out very long runs to guarantee that the evolution dynamics (genetic algorithm) leads the population to an optimum or quasi-optimum communication code, in the sense that the code maximizes the number of successful communication events per agent. We find that, whenever it is possible, the dynamics leads to an ideal code, i.e., a one-to-one mapping between signs and situations.*

1. INTRODUCTION

The uniqueness of human language is probably one of the few, if not the sole, scientific ideas that still resist the corrosive effects of the, borrowing Dennett's metaphor [1], "universal acid" that stems from Darwin's concept of evolution through natural selection. The notion of a "language organ" exclusive of the human species which was originally designed to carry out combinatorial calculations [2] and the exaggerated emphasis on the role of cultural evolution, in opposition to genetic evolution, on the development of language [3,4] are often invoked to support the claim that we are the only species capable of genuine symbolic thinking and communication [5]. This anthropocentric view is usually criticized by ethologists [6,7] who seek to demonstrate that the gap between human and non-human languages is not that big and it is actually magnified by our ignorance about the basic elements used in the communication of non-human animals [7]. Nonetheless, up to now the ethologists have failed to provide clear evidence of, say, syntax in non-human languages. In fact, those languages are typically non-syntactic, i.e., signals refer to whole situations, in contrast to human language which is characterized by signals formed by discrete components that have their own meaning. Together with the language organ, that

composition allows us to take advantage of combinatorics and so as linguists put that to "make infinite use of finite means".

In the 1990s, the ethological approach to the evolution of communication received a rather unexpected ally, namely, computer simulations of large communities of simple finite-state machines endowed with the capacity to emit as well as to respond to signals. This *in silico* approach, termed synthetic ethology by its founder Bruce MacLennan [8], aimed at realizing experiments on the evolution of communication in completely controlled and transparent set-ups, a goal much beyond the empirical capabilities of contemporary ethology.

Before proceeding, it is necessary to provide a working definition of communication. There are almost as many such definitions as authors that have written on the topic of communication (see page 7 of ref. [7] for a sample) but here we follow MacLennan and use Burghardt's definition [9]:

Communication is the phenomenon of one organism producing a signal that, when responded to by another organism, confers some advantage (or the statistical probability of it) to the signaler or his group.

Actually we will assume, as done also by MacLennan, that correct communication about events provides a fitness advantage to *both* signaler and receiver. In this contribution we modify slightly the original synthetic ethology framework introducing independent modules (genes) for the emission of signals and for the actions elicited by those signals. More importantly, we show that earlier criticism and suspicions that a community of agents would not be capable to develop an ideal code, i.e., a one-to-one mapping between signs and situations are unfounded [10].

Use of words signs and symbols in literature is inconsistent. As Deacon noted, symbol is one of the most misused words [5]. In mathematical literature, they are used interchangeably. In semiotic literature usage is inconsistent [11]. In general culture, symbols are

understood as having profound meanings. In analytical Jungian psychology, symbols, are psychological processes connecting conscious and unconscious [12]. In Pribram [13], symbols as adaptive, context-sensitive signals in the brain, whereas signs he identified with less adaptive and relatively context-insensitive neural signals. According to general culture and [5, 12, 13, 14], we use the word sign for notations with predefined meanings, and we reserve the word symbol for psychological processes in which meanings emerge.

In the next section, a variant of the model proposed by MacLennan is introduced and the genetic algorithm governing the evolution of the population of agents is described. Section 3 then presents the results of this model for different values of the sizes of the repertoires of signs and situations. Finally, section outlines the direction of future research.

2. THE MODEL

The model we use in this contribution is a variant of the model proposed in the seminal work of MacLennan [8]. There are N agents that interact among themselves by perceiving and making changes in the environment they inhabit. The environment of each agent is composed of two parts – a public or global environment which is shared with all other agents and a private or local environment, which no other agents have access to. The architecture of the agents' world is illustrated in figure 1. The public environment can be found in a finite number of states, each state represented by the integer $\gamma \in \{1, 2, \dots, G\}$. Similarly, the state of local environment is described by the integer variable $\lambda \in \{1, 2, \dots, L\}$. The basic idea is to permit the agents to exchange information about their local environments by reading and writing (i.e., modifying the state) on the public environment. In that sense, we refer to the state of the private environment as *situation* λ and the state of the public environment as *sign* γ . The goal is then to let the population evolve a mapping between situations and signs.

To accomplish that goal the agents must be endowed with two capabilities (cognitive and motor prerequisites). First, they must be sensitive to the states of those environments, which are actually modeled as input signals to the agents' sensorial channels. Second, the agents must be able to respond actively to the stimuli from the environments. That response is of two types, an *action* and an *emission*. When performing an emission, the agent draws a sign γ' from the set of signs used to describe the public environment and replaces the current state of that environment by γ' . In doing so, the agent modifies the public environment and so emitting is like signalling. On the other hand, an action is a more introspective business: prompted by the sign placed in the public environment another agent draws a situation λ' from the private

environment repertoire. In other words, the agent interprets the signal γ' as meaning situation λ' . Of course, the correctness of this inference will depend on whether the private environment of the agent that last modified the public environment (by writing the sign γ' on it) was λ' . In that case, we say that a successful communication event has occurred and both agents involved in that event are rewarded in the sense of having their fitness increased by one unity of fitness.




Figure 1- The structure of the world inhabited by four agents. There are four private environments

We assume that, once prompted to respond, each agent performs an action and subsequently an emission. This differs from MacLennan's model, in which an agent can either act or emit. In summary, the agents are modeled by finite state machines with two transition tables for each agent

$$A(\gamma, \lambda) \mapsto \lambda' \quad \text{and} \quad E(\gamma, \lambda) \mapsto \gamma' \quad (1)$$

depending on whether the agent is acting or emitting, respectively. In the original implementation the agents act deterministically, i.e., given the inputs γ and λ the agents respond according to (1). Of course, since in principle each agent has a different transition table, the same input can elicit distinct response in different agents. To make things interesting, the private environment of the agents (in other words, its state λ) must change randomly at certain times so that all agents can have access to the entire repertoire of situations. This procedure is necessary since it allows an agent to make use of its entire transition table during its lifetime. In addition, if the private environment were kept fixed then it would be impossible for any single agent to develop a one-to-one mapping between signals and situations as it would have experienced only one situation during its lifetime. Hence

we note that the agent “identity” is its emission and action transition tables, not the state of its private environment. More specifically, let us assume that at each unit of time, which we term an “hour”, in average all agents are prompted to respond to the current stimuli provided by their environments. In other words, in the interval of one hour we choose randomly N agents, one by one, and prompt them to respond to their stimuli. After H hours, an interval of time we call a “day”, the private environments of all agents are modified by choosing randomly situations from the repertoire $\{1, 2, \dots, L\}$. After D days, which comprise the interval of time termed a “week”, we compute the number of successful communication events each agent has participated in and used this quantity as a measure of the fitness of the agent. Then we choose a single agent with probability proportional to its relative fitness (i.e., the agent’s fitness divided by the sum of the fitness of all agents). A copy (clone) of this selected agent is made and some small changes (mutations) are performed on the clone with probability u . More pointedly, we choose randomly a pair sign-situation (γ, λ) and modify, also randomly, the corresponding outputs λ' and γ' of the transition tables for this single pair [see equation (1)]. Finally, in order to keep the population size constant we eliminate the agent with the lowest fitness value. This procedure differs from the standard genetic algorithm implementation [15] in that it allows for the overlapping of generations, a crucial prerequisite for cultural evolution which may be relevant in the case when learning is allowed.

Before proceeding with the presentation of the simulation results of this minimal model for the evolution of communication, we note that only a few control parameters can significantly affect those results. In fact, the whole issue boils down to finding the structure of the transition tables that maximizes the number of successful communication events in a population composed of N agents, whose behavior is determined by those tables. In this perspective, the genetic algorithm is simply a means to find that optimum and hence the choice of parameters u , H , D , as well as of the duration of the run in unit of weeks, which we call W , the mode of reproduction (presence or absence of crossover and overlapping or non-overlapping of generations) can affect our ability to reach the maximum, but not the properties of the maximum itself. Those properties depend only on the parameters G and L , the sizes of the repertoires of signs and situations, respectively. Hence, in what follows we will solely present results for the set of parameters that produced the best communication accuracy. (Of course, the ideal communication corresponds to $A(\gamma, \lambda) = \lambda$, $E(\gamma, \lambda) = \gamma$, for all agents; which is only possible if $G=L$).

3. RESULTS




Figure 2 Best and average fraction of successful communication events as function of the number of weeks. The parameters are $N = 100$, $H = 10$, $D = 5$, $u = 0.1$ and $G = L = 8$. The lines at the bottom are the best and average results for chance guessing.

The relevant quantity to study is clearly the average number of successful communication events per agent, since this is the measure one seeks to maximize. Also important is the number of successful communication events of the fittest agent in the population. In figure 2 we plot these quantities as function of the time measured in units of weeks. The total simulation time was 4×10^6 , about three orders of magnitude greater than the typical runs performed in the original experiments [8]. In fact, MacLennan’s analysis focused mainly on the rate of increase of the mean fitness of the population, calculated through a linear fitting of the smoothed data. This approach, however, makes little sense nowadays when computer resources allow us to carry out much longer runs. For the purpose of comparison, we also present in figure 2 the results for the case that communication is suppressed, i.e., the only possibility of successful communication event is purely by chance (“guessing”). This is achieved by writing a random sign at the public environment instead of the sign encoded in the transition table $E(\gamma, \lambda)$. We note that any successful communication event, regardless of whether achieved by pure chance or through adaptation, is rewarded. Interestingly, at the end of the run (with communication enabled) about 90% of the communication event are successful – this is well above the chance level values of 12,5% for the average and 20% for the best performances. Moreover, MacLennan found that the fitness seemed to mysteriously increase, although extremely slowly, with time even in the case communication is suppressed: the two horizontal lines in figure 2 depicting the best and average performances of the guessing strategy demonstrate that

this spurious effect does not appear in the present set up where emission and action are considered separately.

To better understand the communication code evolved by the population of agents we should look at its *denotation matrix*, the elements of which $D_{\gamma\lambda}$ yield the fraction of times a pair sign-situation (γ, λ) is used successfully in a certain number of communication events. The denotation matrix is computed for successful events only since, at this stage, we assume that the agents have developed a communication code and successful communication is the result of using that code. In particular, considering the last 1844308 successful communication events of the run described in figure 2 we find that the only non-vanishing elements of the denotation matrix are $D_{18} = 0.103$, $D_{21} = 0.100$, $D_{36} = 0.120$, $D_{43} = 0.136$, $D_{57} = 0.134$, $D_{64} = 0.135$, $D_{72} = 0.134$ and $D_{85} = 0.137$. This result indicates that the agents managed to evolve a one-to-one correspondence between signs and situations – an ideal communication code. Of course, any permutation of this code yields an equally optimal solution. More importantly, perhaps, this result dispels the suspicion that in seeking for an optimal communication code the agents would tend to decrease their repertoire of signs [10]: inspection of the entries of the denotation matrix indicates that all signs are used with approximately equal frequencies. The reason that the repertoire of signs is not decreased is that our model rewards the differentiated understanding and communication about the environment: agents will attempt to use as many communication signs $\{\gamma\}$ as there are situations $\{\lambda\}$. The highly structured denotation matrix contrasts with the practically uniform values of the entries of the denotation matrix in the case communication is suppressed (data not shown).

These findings encourage us to proceed to a closer examination of the transition tables of the agents that survived at the end of the run. In fact, we find that those agents share in average 85% of the entries of the transition tables, i.e., the surviving agents are practically identical. Actually, what prevents the population of becoming completely homogeneous is the diversity introduced by the mutations during the copying process. Examination of the transition tables of the best communicator revealed the secret of its success: for the pairs sign-situation (γ, λ) for which the entries of the denotation matrix are non-zero we find $A(\gamma, \lambda) \mapsto \lambda$ and $E(\gamma, \lambda) \mapsto \gamma$, i.e., the agent can communicate perfectly with itself or with any of its non-corrupted clones. From the evolutionary biology viewpoint this kind of result is not surprising, since coexistence of distinct replicator species is very difficult to achieve and necessitates a special selection pressure to favor it, namely, group selection [16]. Perhaps, related to this finding is Chomsky's notion of a Universal Grammar

that provides the foundation to all human languages (see, e.g., [17]).




Figure 3 - Average fraction of successful communication events as function of the number of weeks. The parameters are the same as those of figure 2 except for G that takes on the values ($\leq L = 8$) shown in the figure.

Once demonstrated the suitability of our framework to study the evolution of communication codes among agents modeled by finite-state machines, we consider now the more general case, in which the sizes of the repertoires of signs and situations differ, i.e., $G \neq L$. Let us consider first the case in which there are more situations than signs to express them ($L \geq G$). Figure 3 illustrates how the average fraction of successes in communication events evolves with time (in weeks) using the same parameters of the genetic algorithm as before, but with G varying from 2 to 8 while L is kept fixed at $L = 8$. We recall that the average performance of the random guessing strategy is $1/L = 0.125$ regardless of the value of G .

Inspection of the denotation matrices and the transition tables of the best communicators indicate that the genetic algorithm has found the optimal solution in each case (average fraction of successes approximately equal to G/L). Moreover, although a one-to-one assignment between signs and situations is now impossible, we have verified that each situation is assigned to only one sign (of course, this sign may be used to express other situations as well). We note that only the emission strategy must be finely tuned in this setting. Consider, for instance, the extreme case $G = 2$ and $L = 8$. By reading the sign displayed in the public environment, an agent has four distinct options of action – all of them successful. However, once it has performed an action there is only one option for emission to match that action. We can

actually see the effect of these constraints in the structure of the transition tables of the agents at the end of the run: they share 81% of the entries of emission table $E(\gamma, \lambda)$, but only 43% of the entries of the action table $A(\gamma, \lambda)$. In other words, selection is strong for the emission part of the agent's genome, but weak for the action part.

A much easier problem from the optimization perspective is the case that there are more signs than situations to be described, $L \leq G$. Figure 4 shows the time evolution of the fraction of successful communication events averaged over all agents in the population.




Figure 4- Average fraction of successful communication events as function of the number of weeks. The parameters are the same as those of figure 2 except for L that takes on the values ($\leq G = 8$) shown in the figure.

The strategy found by the genetic algorithm was to discard the surplus of signs and then to produce a one-to-one correspondence between the remaining signs and the situations, using the same optimal scheme discovered in the case $L = G$. Since the entries of the transition tables involving the discarded signs are not used, we should expect a great diversity among the agents, as far as those entries are regarded. In fact, this is what we generally found when comparing the entries of the transition tables of all agents at the end of the run. For instance, in the run depicted in figure 4 for $G=2$ we find that the agents share only 31% of the action and 25% of the emission entries. There are, however, other types of equally optimum solutions that were found in different runs and leads to a completely different composition of the population. For example, a run with $G=8$ and $L=2$ resulted in a denotation matrix that assigns signs $\gamma = 2, 4$ and 6 to situation $\lambda = 1$. Inspection of the genome of the best communicator revealed the strategy $A(2,1) = A(4,1) = A(6,1) = 1$ for action and $E(2,1) = 6$,

$E(4,1) = 2$, $E(6,1) = 4$ for emission. There is clearly more freedom in choosing the emission strategy (any permutation of the signs 2, 4, 6 will be equally good) than the action strategy. In fact, in that run we found that the agents shared 62% and 18% of the entries of the action and emission, respectively, transition tables.

4. CONCLUSION

Given the spaces of meanings and signals, and a notion of success in a communication event we should foster no doubts that the emulation of evolution by natural selection brought about by the genetic algorithm will produce an optimum communication code among the agents. The next challenge is to adapt the present framework to study the emergence of compositional or syntactic communication codes. Up to now studies of the evolution of syntactic communication have either assumed the existence of such codes and then focused on the conditions for natural selection to favor syntactic over non-syntactic codes [18] or employed sophisticated algorithms to produce the rules of the grammar [19]. Interestingly, a simple modification of our variant of MacLennan's model may suffice to produce syntactic communication codes: allowing a variable number of signs to be displayed simultaneously in the public environment and considering a repertoire of situations that is much larger than the repertoire of signs, though much smaller than the number of their combinations, may lead to the emergence of compositional codes. Work in this direction is on the way.

ACKNOWLEDGMENTS

Effort sponsored by the Air Force Office of Scientific Research, Air Force Material Command, USAF, under grant number FA8655-04-1-3045. The U. S Government is authorized to reproduce and distribute reprints for Governmental purpose notwithstanding any copyright notation thereon.

REFERENCES

- [1] D. C. Dennett, *Darwin's Dangerous Idea*, New York: Simon & Schuster, 1995.
- [2] J. Fodor, *The Modularity of Mind*, Cambridge: MIT Press, 1983.
- [3] S. Blackmore, *The Meme Machine*, Oxford: Oxford University Press, 1999.

- [4] K. Smith, S. Kirby and H. Brighton, "Iterated Learning: a framework for the emergence of language," *Artificial Life* **9**, 371-386, 2003.
- [5] T. W. Deacon, *The Symbolic Species*, New York: W.W. Norton & Company, 1997.
- [6] R. Dunbar, *Grooming, Gossip, and the Evolution of Language*, Cambridge: Harvard University Press, 1996.
- [7] M. D. Hauser, *The Evolution of Communication*, Cambridge: MIT Press, 1996.
- [8] B. J. MacLennan, "Synthetic ethology: an approach to the study of communication," *Artificial Life II*, SFI Studies in the Sciences of Complexity, vol. X, 631-658, Addison-Wesley, 1991.
- [9] G. M. Burghardt, "Defining Communication," in *Communication by Chemical Signals*, edited by J. W. Johnston Jr., D. G. Moulton and A. Turk, New York: Appleton-Century-Crofts, 1970.
- [10] J. Noble and D. Cliff, "On simulating the evolution of communication," *Proceedings of the 4th International Conference on Simulation of Adaptive Behavior*, Cambridge, MA, MIT Press, 1996.
- [11] T. A. Sebeok, *Sign: An Introduction to Semiotics*. Univ. Toronto Press, Toronto, Canada. Peirce, C.S. (1935-66). *Collected Papers of Charles Sanders Peirce*. Harvard Univ. Press., Cambridge, MA; Morris C. (1971). *Writings on the General Theory of Signs*, ed. Th. A. Sebeok. The Hague: Mouton, 1995.
- [12] C. G. Jung, *Archetypes of the Collective Unconscious*. In the *Collected Works*, v.9,II, Bollingen Series XX, 1969. Princeton: Princeton University Press, 1934.
- [13] K. Pribram, *Languages of the Brain*. New York: Prentice Hall, 1971.
- [14] L. I. Perlovsky, *Neural networks and intellect: using model based concepts*. New York: Oxford University Press, 2001.
- [15] M. Mitchell, *An Introduction to Genetic Algorithms*, Cambridge: MIT Press, 1996.
- [16] S. A. Boorman and P.R. Levitt, *The Genetics of Altruism*, New York: Academic Press, 1980.
- [17] M. A. Nowak, N. L. Komarova and P. Niyogi, "Evolution of Universal Grammar," *Science* **291**, 114-118, 2001.
- [18] M. A. Nowak, J. B. Plotkin and V. A. A. Jansen, "The Evolution of Syntactic Communication," *Nature* **404**, 495-498, 2000.
- [19] S. Kirby, "Spontaneous Evolution of Linguistic Structure: an iterated learning model of the emergence of regularity and irregularity," *IEEE Trans. Evol. Comput.* **5**, 102-110, 2001.

Minimal models for text production and Zipf's law

José F. Fontanari

Universidade de São Paulo, São Carlos, Brazil, fontanari@ifsc.usp.br

Leonid I. Perlovsky

Air Force Research Lab, Hanscom AFB, USA, Leonid.Perlovsky@hanscom.af.mil

Abstract — *Mathematical models used to explain the power-law distribution of word frequencies observed in natural languages - Zipf's law – generally assume that symbols and words occur independently, i.e., they do not interact. Here we show that when interaction is taken into account by allowing the words to compete amongst themselves for space in the memory of the users, the resulting word frequency distribution is best described by an exponential, rather than by a power-law. The implications of the failure to derive Zipf's law under more realistic assumptions are discussed.*

1. INTRODUCTION

The notion that words compete and languages evolve similarly to individuals and populations was already familiar in Darwin's time. In fact, the following quote from Darwin makes the point clear [1]:

We see variability in every tongue, and new words are continually cropping up; but as there is a limit to the powers of the memory, single words, like whole languages, gradually become extinct. As Max Müller has well remarked: - "A struggle for life is constantly going on amongst the words and grammatical forms in each language. The better, the shorter, the easier forms are constantly gaining the upper hand, and they owe their success to their own inherent virtue." To these more important causes of the survival of certain words, mere novelty may, I think, be added; for there is in the mind of man a strong love for slight changes in all things.

We refer the reader to Ref. [2] for a detailed account of Darwin's contribution to the debate on language change as a selection process. More recently, the well-documented development of Romance languages from Latin (i.e., the gradual divergence of the languages of

France, Italy, Spain, Portugal and Romania from Latin, as well as from each other) has offered a convincing proof that groups of related languages develop and diverge from a common ancestral tongue, similarly to gene lineages [3]. In view of these observations, one should not be surprised to encounter evolutionary arguments and population-genetics-inspired mathematical models playing a leading role in the explanation of features of language. In this contribution, we use this evolutionary approach in seeking to comprehend a quite remarkable aspect of natural (i.e., produced by humans) texts, namely, Zipf's law [4].

In the early 1930s George Zipf noticed that if a large sample of words in a text are arranged in rank order, from most frequent to least frequent, then the dependence of the frequency f of a word on its rank r is very well described by the power-law distribution $f \propto 1/r$, regardless of the language or speaker [4]. The significance of Zipf's law in language, however, is still an unsettled issue. On the one side, some authors, arguing that texts produced by the random emission of symbols also generate word frequency distributions that follow Zipf's law (more precisely, the generalized Zipf's law), claim that this law is linguistically very shallow [5,6]. On the other side, some authors point out that the fact that random systems display Zipf-law-like distributions does not exclude the possibility of Zipf's law being a genuine reflex of mechanisms underlying the behavior of complex systems [7]. In other words, it is argued that the random emission of symbols is simply not a valid null model for the creation of texts in natural languages [8,9]. A valid model should be based on realistic assumptions on the factors that originate natural texts. Following the suggestion of the renowned philologist of the 19th century Friedrich Max Müller mentioned in Darwin's quotation, our guide in this endeavor will be the theory of evolution by natural selection.

In the next section, we describe a branching evolutionary model that results in word frequency distributions that obey Zipf's law and then propose a change in that model in order to take interactions between words into account. The rank statistics of this variant is then investigated in section 3. Finally, section 4 summarizes our main results and indicates directions for future research.

2. THE MODEL

We begin by reviewing a simple evolutionary model that produces a non-stationary distribution of word frequencies that obeys Zipf's law [10, 11]. That model is usually formulated in the language of ecological dynamics (i.e., the basic elements are individuals that are categorized in different species), but here we will face the challenge of presenting the model solely in linguistic terms. In particular, we will term *word store* the linguistic analogous of an ecosystem.

At any given time t the word store is completely characterized by the set of integers $n_k(t)$, $k = 1, 2, \dots, K(t)$, where $n_k(t)$ is the number of times word k appears in the word store and $K(t)$ is the size of the vocabulary (i.e., the number of different words in the word store). We assume time is discrete and increases in steps of unitary size. At each step exactly one word is created: it can be a new word and so increase the vocabulary size by one or a copy of a word already present in the word-store. The probability that a new word crops up at step $t + 1$ is defined as

$$\Pr[n_{K+1} = 1 | n_{K+1} = 0] = c \quad (1)$$

where $c \in [0, 1]$ can be viewed as the mutation probability. The probability that a known word k is created is

$$\Pr[n_k + 1 | n_k] = (1 - c) \frac{n_k}{N} \quad (2)$$

where $N = \sum_{k=1}^K n_k$ is the total number of words in the word store and, for simplicity, we have omitted the dependence on t . The conditions at $t = 0$ are fixed as $n_1(0) = 1$ and $K(0) = 1$. Equation (2) indicates that the more frequent a word is, the more frequent it will become, which is essentially the basic assumption of the so-called discourse-triggered word choice model [8, 9]. In addition, this is also the usual assumption used in population genetics to model neutral evolution [12].

Before considering changes in this standard model, let us illustrate some of its predictions. For instance, in figure 1 we show the dependence of the frequency n_k/N on the re-scaled rank $r/\langle K \rangle$ where $\langle K \rangle = ct$ is the average vocabulary size for runs of duration $t_m = 4 \times 10^4$ and different values of c . The results are averages over 1000 independent runs. Increasing the duration of the runs does not affect the results exhibited in the figure. The straight line in a double logarithmic scale is the signature of Zipf's law so the model is quite successful in predicting this feature of language.




Figure 1 Frequency against rank re-scaled by the average vocabulary size for (top to bottom) $c = 0.01$, 0.1 and 0.5 . The lines are the linear fittings which have slopes -1.11 , -0.96 and -0.58 respectively.

There are at least two points of departure between the evolutionary model just presented and Darwin's view of language evolution. First, the words do not become extinct, which is also in disagreement with estimates from glottochronology (i.e., the chronology of languages) that suggest the rule of thumb that languages replace about 20 percent of their basic vocabulary every one thousand years [3]. Second, there is actually no competition or "struggle for life" among words. In fact, the very reason for branching Markov processes being amenable to analytical approaches is because there is no interaction between branches, i.e., they evolve independently of each other. In order to address these two points while keeping most basic features of the evolutionary model unaltered we add a stochastic death process after the birth of a word has taken place, regardless whether according to process (1) or (2). Explicitly, after the birth of a word we pick randomly a word from the word-store of size N and eliminate it with probability

$$P_{\text{death}} = \exp[-\beta(M - N)] \quad (3)$$

if $N \leq M$ and $P_{\text{death}} = 1$, otherwise. Here M is the carrying capacity of our "memory" and the smoothness β is a free adjustable parameter. This modification will not change the dynamics in the initial steps ($N \ll M$), but in the asymptotic regime it will lead to a saturation of the size of the word store, exactly as implemented in the classic Moran model of population genetics [12]. The mechanism to keep this size fixed introduces thus an effective competition between words. To clarify this important point, let us consider the case $c = 0$. If we start the run of the branching process with two distinct words, i.e., $n_1(0) = n_2(0) = 1$ and so $K(0) = 2$, we will always find these two words in the word store. In the competition model, however, passed some time only one

of the words will be found in the word store. In absence of mutations, competition leads ultimately to the dominance of a single type of word.

It is interesting to note that in the model of language evolution investigated here, which is inspired in Darwin's and Müller remarks quoted in the beginning of the paper, words compete for space in the memory of the language users. Nowhere is it said that words confer fitness to those users who then compete among themselves. Hence our model is one of *memetic*, rather than genetic, evolution (see, e.g., [13]).

3. RESULTS

In what follows we set the smoothness parameter to $\beta = 1$ and leave the word store to evolve until $t_m = 8 \times 10^4$. At this point the quantities of interest are measures and stored for statistical purpose. The data presented in the next figures are averages over 5000 runs. First we note that the vocabulary size K cannot increase linearly with time as in the original model, since it is obviously bounded by the carrying capacity M . In figure 2 we show the dependence of the ratio K/M on c for $M = 2000$ at the stationary regime. This result is not affected by different choices of the memory capacity M , indicating thus that $K \propto M$.




Figure 2 Ratio of vocabulary size to carrying capacity of the word-store as a function of the mutation probability in the stationary regime.

We turn now to the analysis of the rank statistics. In figure 3 we present an analogue of figure 1 for the model with competition. For the sake of comparison, the rank is re-scaled by $\theta = Mc$. In fact, since in the original model a word is created at each time step, so that the run time t equals the word store size M , it is clear that θ is

equivalent to $\langle K \rangle$. More important, this re-scaling becomes identical to that used in ref. [9], when one takes into account the factor 2 in the definitions of θ for the Moran model used here and the Wright-Fisher model used in ref. [9] (we refer the reader to the book by Ewens [12] for the explanation of this subtle point). The results are presented in figure 3 using a semi-logarithmic scale so that fitting by a straight line indicates an exponential rather than a power-law frequency distribution. In fact, for small c the exponential yields the best fitting, in agreement with the analytical predictions of ref. [9] but in disagreement with the preliminary numerical results of ref. [8]. We note that there is an intrinsic difficulty to produce a representative frequency distribution for, say $c = 0.5$ since according to figure 2 about half of the words in the word-store are different and so the degeneracy n_k of each word is simply too small to validate the rank statistics.




Figure 3 Frequency against rescaled rank in a semi-logarithmic scale for (top to bottom) $c = 0.01$, 0.1 and 0.5 . The line is the exponential fitting for the lower mutation probability and yields the slope -1.15 .

These technical difficulties are absent in the analytical approach of ref. [9] because M and K are made arbitrarily large.

4. CONCLUSION

In looking for a motivation to introduce the branching evolutionary model in the linguistic context, Günther et al offered the reader a remarkable insight: Zipf's law is usually derived under the assumption of non-interacting particles (interpreted as symbols, words, etc.), analogously to the "ideal gas" of thermodynamics [7]. For instance, in the influential paper by Li on random texts, symbols from an alphabet that includes the blank space are generated independently which is equivalent to assume they do not interact [6]. However, as far as the presence of interactions is concerned, the branching

process evolutionary model [7,10,11] does not differ from the more explicit ideal gas models. As already pointed out, each lineage evolves independently of each other and so that model fails to take interactions into account. In view of the above remarks, the rank analysis of the branching evolutionary model leads to Zipf's law (as illustrated in figure 1) because of independent evolution and no interaction.

In this contribution we have shown that if competition among words, that results from the limited capacity of memory of the language users, is incorporated into the original evolutionary model then the words frequency distribution becomes an exponential rather than a power-law (see figure 3): Zipf's law is not recovered. This is a most interesting finding because it implies that either there is no such a thing as a "struggle for life" amongst words and so they evolve independently or then the concept of evolution through natural selection is not suitable to describe the evolution of language. Perhaps culture (see, e.g. [14]) is the missing ingredient needed to derive Zipf's law under the more realistic assumption of interaction among words.

ACKNOWLEDGMENTS

Effort sponsored by the Air Force Office of Scientific Research, Air Force Material Command, USAF, under grant number FA8655-04-1-3045. The U. S Government is authorized to reproduce and distribute reprints for Governmental purpose notwithstanding any copyright notation thereon.

REFERENCES

[1] C. Darwin, *The Descent of Man, and Selection in Relation to Sex*, vol. 1, pp. 60-61. Facsimile edition, eds. J. T. Bonner and R. M. May, Princeton : Princeton University Press, 1981.

[2] G. Radick, "Darwin on Language and Selection," *Selection* **3**, 7-16, 2002.

[3] C. Renfrew, *Archaeology and Language*, Cambridge: Cambridge University Press, 1987.

[4] G. K. Zipf, *Human Behavior and the Principle of Least Effort*, Cambridge: Addison-Wesley, 1949.

[5] B. Mandelbrot, *The Fractal Geometry of Nature*, New York: Freeman, , 1982.

[6] W. Li, "Random texts exhibit Zipf's-law-like word frequency distributions," *IEEE Trans. Inf. Theor.* **38**, 1842-1845, 1992.

[7] R. Günther, L. Levitin, B. Schapiro and P. Wagner, "Zipf's Law and the Effect of Ranking on Probability Distributions," *Int. J. Theor. Phys.* **35**, 395-417, 1996.

[8] C. Tullo and J.R. Hurford, "Modelling Zipfian Distributions in Language," *Proceedings of the Language, Evolution and Computation Workshop/Course* at the 15th European Summer School on Logic, Language and Information, Vienna, 2003.

[9] J. F. Fontanari and L. I. Perlovsky, "Solvable null model for the distribution of word frequencies," *Phys. Rev. E* **70**, 042901, 2004.

[10] B. Schapiro, "An approach to physics of complexity," *Chaos, Solitons and Fractals* **4**, 115-123, 1994.

[11] L. Levitin, B. Schapiro and L.I. Perlovsky, "Zipf's Law Revisited: Evolutionary Model of Emergent Multi-resolution Classification," *Proceedings of the Conference on Intelligent Systems and Semiotics '96*. Gaithersburg, MD, v.1, pp. 65-70, 1996.

[12] W. J. Ewens, *Mathematical Population Genetics*, 2nd edition, New York: Springer, 2004.

[13] S. Blackmore, *The Meme Machine*, Oxford: Oxford University Press, 1999.

[14] R. Boyd and P. J. Richerson, *Culture and the Evolutionary Process*, Chicago: University of Chicago Press, 1985.

Meaning Creation and Modeling Field Theory

José F. Fontanari

Universidade de São Paulo, São Carlos, Brazil, fontanari@ifsc.usp.br

Leonid I. Perlovsky

Air Force Research Lab, Hanscom AFB, USA, Leonid.Perlovsky@hanscom.af.mil

Abstract — We study the development of the discriminatory capacity (i.e., the ability to develop a concept or categorize each object in the environment) of a single organism using two distinct approaches, namely, discriminatory trees and Modeling Field Theory (MFT). In particular, we consider a simple world composed of objects that are characterized by real-valued features, similar to that used in seminal works on meaning creation. Within that framework, we demonstrate in a series of didactic experiments the potential of the MFT approach as a truly autonomous (as opposed to discrimination trees) mechanism for meaning generation

1. INTRODUCTION

A major criticism against traditional agent-based models of language evolution is that the agents are always provided with *a priori* structured meaning spaces that ultimately are responsible for all observed “emergent” properties (e.g., syntactic structure) of the evolved language [1]. (The word emergent is written between quotation marks to remind us of Minsky’s assertion that the use of the word “emergence” should make one suspicious that not enough effort has been made in finding explanatory mechanisms [2].) In other words, there is no creativity of concepts in those models. The fact that conceptual knowledge is fixed at the outset precludes analysis of more plausible scenarios in which meaning and linguistic representations are generated concurrently, enhancing each other. This alternative framework for computer modeling language or communication evolution was put forward by Steels [3] and explored further by Smith [4].

To instantiate any model of communication between virtual or real organisms, a basic cognitive requirement must be fulfilled, namely, that the organisms be capable of classifying different types of situations and, accordingly, be capable of recognizing that a situation of a particular type turns up. In this vein and for the purpose of this paper, meaning is viewed as a categorization of reality which is relevant from the perspective of the

organism. Hence meaning creation is synonymous to category creation, i.e., the ability to distinguish, through the creation of internal representations or concepts, the objects, as well as the other organisms, that make up the organism’s *Umwelt* (ethologist’s jargon for the environment in which an organism is embodied and embedded). This is achieved through a generalization of Wittgenstein’s notion of language games [5] to the non-linguistic domain, resulting in the so-called *discrimination* games [3]. In these games an organism inhabits a simple world made up of N objects or situations, each of which is described by a single feature value modeled by a real variable $O_i \in (0,1), i = 1, \dots, N$ drawn randomly from a uniform distribution. These features are, of course, abstract and have no particular meaning in the model, though it may be helpful to think of them as perceptual features such as color or smell. The question is whether such organism is able to form *autonomously* a repertoire of features to succeed in discrimination and to adapt that repertoire when new objects are considered. In this contribution we address this problem using both the original discrimination tree approach [3] and a novel adaptive approach to concept formation dubbed modeling field theory [6].

Following Steel’s original paper, we will consider meaning creation in a single agent, so that the communication issue is not addressed at this stage (see [4] for the natural extension of this research program to study communication in a community of agents). However, rather than considering that each object is characterized by a set of features and that each organism has a set of sensory channels designed to detect each feature (there is a one-to-one mapping between channels and features), here we assume that there is only one feature per object and that the organisms possess a single sensory channel sensitive to that feature value. Creation of meanings in high-dimensional spaces, as well as extending a notion of object to abstract concepts will be a subject of future publications.

In the next section, we review the approach of discrimination trees to meaning creation. In particular, quantitative performance measures are presented for both the standard algorithm in which refinements of the tree are undertaken randomly and the intelligent tree growth strategy in which a refinement always make a helpful distinction. In section 3 we briefly review the modeling field theory approach and describe the results of its

application to the categorization problem posed above. Finally, section 4 summarizes the main conclusions.

2. DISCRIMINATION TREES

The idea of the discrimination trees is to model the sensitive channel by a binary tree as illustrated in figure 1. The nodes of this tree are labeled unambiguously by a binary sequence (e.g., 010) and are endowed with the capacity to detect whether a feature value falls between two bounds, except for the root (node 0) that has no discriminatory power - it is sensitive to the entire range 0-1. Meaning creation takes place by splitting the sensitivity range of a node in two, resulting thus in the production of two new nodes, each one sensitive to half of the range of values of the parent node. Hence, for example, node 00 is sensitive to features whose values are within the range 0 - 0.5; node 01 to values within the range 0.5-1 and node 0100 to values within the range 0.5-0.625. The sensitive channel represented by the tree shown in the figure is capable to distinguish between, say, objects $O_i = 0.6$ and $O_j = 0.7$, but fails to distinguish between objects $O_k = 0.1$ and $O_l = 0.4$. The final discrimination capability of the tree is determined by its *leaves* (i.e., the external nodes 00, 011, 0101 and 0100 in the example). In fact, to perfectly categorize N objects a tree must possess at least N leaves. It is also useful to define the *depth* of a node as the minimum number of branches connecting it to the root and the depth of a tree as the maximum of the depth of its nodes.




Figure 1 – Discrimination tree with four leaves, three internal nodes and depth equal three. This tree is sensitive to features values in the ranges (0,0.5), (0.5,0.625), (0.625,0.75) and (0.75,1).

In this publication we follow [3] in assuming that Umwelt is populated by objects, and the meaning creation consists in learning to differentiate them. We do not address the issue of learning abstract concepts, which are not represented by individual objects. In this context it is the failure to distinguish between any two objects that leads to further splitting or refinement of the discrimination tree and hence to improvement of the semantic structure of the

sensory channel. This is done through repeated discrimination games, in which one of the N objects that compose the organism's world is chosen randomly and compared with the $N-1$ remaining objects. Whenever a failure occurs a particular leaf is split into two new leaves, creating thus a pair of (derived) novel concepts in the semantic structure of the channel. (We would like to emphasize that the assumption of the finite number of distinct objects is a significant simplification of the current publication as well as of [3], still a step toward complexity of creating novel concepts in the real world, as compared to the current state of the art, e.g. [1]). In what follows we will consider two strategies for the refinement of the tree, namely, the random refinement and the intelligent tree growth.

As the name indicates, in the random refinement strategy, that was used by Steel in his analysis, one chooses randomly, i.e., with equal probability, any of the leaves of the tree and then split it. In the example of figure 1, this amounts to pick randomly one of the four leaves 00, 011, 0101 or 0100. Suppose leaf 011 is chosen. Then the new leaves 0110 and 0111 are created and the parent 011 becomes an internal node. The refined tree has now five leaves and four internal nodes. This example exposes a drawback of the random refinement strategy: although the main shortcoming of the depicted discrimination tree is clearly the failure to distinguish between objects characterized by feature values in the range (0,0.5), the splitting of the node 011 only makes the situation worse – the odds of picking leaf 00 has now decreased from 1/4 to 1/5. Hence an unbalanced tree tends to become even more unbalanced. [We note that ultimately node 00 will be chosen since the probability that it is not chosen in m refinements is $1/(m+1)$.] These remarks are necessary to emphasize that even for a relatively large, though finite, number of discrimination games, the random strategy may fail to create a unique meaning (i.e., leaf) for each object.

Two interesting measures to evaluate the performance of the different strategies are the average number of leaves and the average depth of the discrimination trees. In figure 2 we show these quantities for the random refinement case. More pointedly, we generate 4×10^3 realizations of the N objects by drawing random numbers from the uniform distribution in (0,1). For each realization we repeat the discrimination games 10^4 times or until a perfect categorization of the N objects is achieved. Only these realizations are considered for the evaluation of the averages. Figure 2 summarizes our findings. The fittings indicate that the average number of leaves increases exponentially with the number of objects, $leaves \approx 2.41 \exp(0.71N)$, while the average tree depth increases linearly, $depth \approx 0.15 + 1.14N$. We note, however, that these averages are of little significance since the dispersion around them are very large, especially regarding the number of leaves. For example, in the case

$N = 5$, one of the 4×10^3 instances we used resulted in a tree with 22385 leaves. It is the computer resources needed to keep track of such large trees that limited our analysis of the random refinement strategy to small collections of objects.




Figure 2 – Semi-logarithm plot of the average number of leaves and average depth of the discrimination trees produced by the random refinement strategy against the number of objects N . The lines are the fittings given in the text.

We turn now to the analysis of the intelligent tree growth strategy proposed by Smith [4]. As before, refinement is triggered by a failure in discriminating a given object, say i , from the remaining $N-1$ objects that make up the organism's world. However, in this case one refines the leaf associated to object i , rather than a randomly chosen leaf. For example, consider the tree depicted in figure 1 and assume there are 3 objects with feature values $O_1 = 0.2, O_2 = 0.4, O_3 = 0.7$. If object 3 is chosen to play the discrimination game then nothing happens since leaf 0101 singles out this object from the other two. However, if object 2 is chosen, then a failure occurs because leaf 00 cannot distinguish it from object 1. The procedure is then to refine leaf 00, producing leaves 000 and 001. The latter will provide a unique representation to object 2. This scheme generates optimal discrimination trees, in the sense that the trees possess the minimum number of leaves needed to categorize perfectly the N objects. In contrast to random refinement, the intelligent tree growth strategy produces the same tree for a fixed collection of objects. In figure 3 we show the average number of leaves and the average depth of the discrimination trees produced by this optimal refinement scheme. For fixed N , each data point represents the average over 10^4 realizations of N objects drawn randomly from the uniform distribution. We find that the data for the average number of leaves is very well fitted by the straight line $leaves \approx 1.44N \approx \sqrt{2}N$ while the average depth by the logarithm fitting, $depth \approx 3 \ln N$.




Figure 3 – Average number of leaves and average depth of the discrimination trees produced by the intelligent tree growth strategy as function of the number of object N . The lines are the fittings given in the text.

Perhaps the main criticism one can raise against the discrimination trees approach, and specially regarding the intelligent tree growth strategy, is that one can hardly say that the organism is genuinely autonomous. In fact, the decision of what leaf to refine as well as how to refine it is built in the simulation code. We need a system that is capable to effect that kind of refinement entirely by itself (i.e., autonomously). This is the issue we address in the next section.

3. MODELING FIELD THEORY

The basic idea behind Modeling Field Theory (MFT) is the association between lower-level signals (e.g., inputs) and higher-level concept-models (internal representations) avoiding the combinatorial complexity inherent to such a task. This is achieved by using measures of similarity between concept-models and input signals together with a new type of logic, so-called fuzzy dynamic logic. We refer the reader to Perlovsky's book [6] for a complete presentation of MFT; here we particularize the general framework to the problem of categorizing N objects, each of which characterized by a real number $O_i \in (0,1)$ - the input signals - as described in the previous section. Let us assume that there are M concept-models described by real-valued variables $S_k, k = 1, \dots, M$ that should represent the objects $O_i, i = 1, \dots, N$. We define arbitrarily the following partial similarity measure between object i and concept k

$$l(i | k) = (2\pi\sigma_k)^{-1/2} \exp\left[-(O_i - S_k)^2 / 2\sigma_k^2\right], \quad (1)$$

where, at this stage, the fuzziness σ_k is a parameter given *a priori*. The goal is to find an assignment between models and objects such that the global similarity

$$L = \prod_i \sum_k l(i|k) \quad (2)$$

is maximized. We can easily be deceived by the apparent trivialness of this task, since the categorization mechanisms built in our minds immediately sprout a one-to-one (if $N = M$) correspondence between objects and concepts. However, if asked to formalize that mechanism, the solutions proposed are usually very sophisticated, such as the discrimination trees discussed before. The key point in this task seems to be the symmetry-breaking of the permutation group associated to the labeling of objects by concepts. MFT provides an ingenious method to implement that partition in a fully autonomous framework. A fundamental role is played by the fuzzy association variables $f(k|i)$ defined by

$$f(k|i) = l(i|k) / \sum_{k'} l(i|k') \quad (3)$$

which give a measure of the correspondence between object i and concept k relative to all other concepts k' . A mechanism of concept formation and learning, an internal dynamics of the modeling fields is defined as

$$dS_k/dt = \sum_i f(k|i) [\partial \log l(i|k) / \partial S_k]. \quad (4)$$

It can be shown that this dynamics always converges to a (usually local) maximum of the similarity L . However, by properly adjusting the fuzziness σ_k according to the fuzzy association variables $f(k|i)$ the global maximum can be singled out.

Before considering the full implementation of the MFT scheme, let us first study the dynamics (4) in the case σ_k are fixed. It is important to understand the roles played by the local maxima of L , as (spurious) attractors of the modeling field dynamics. For the sake of concreteness, let us consider five objects ($N = 5$) with features $O_1 = 0.1, O_2 = 0.2, O_3 = 0.3, O_4 = 0.4, O_5 = 0.5$. The number of model-concepts equals the number of objects, i.e., $M = 5$ but to make the task more difficult, the initial values of the modeling fields $S_k(t=0)$ are chosen randomly in the range $(0.5, 1)$. Explicitly, in the experiments reported here we use the following values $S_1 = 0.94, S_2 = 0.59, S_3 = 0.62, S_4 = 0.86, S_5 = 0.79$. The differential equations (4) are solved with Euler's method using the step-size $h = 10^{-4}$. In figure 4 we show the time evolution of the modeling fields when the fuzziness are set to $\sigma_k = 0.15$ (or any value greater than this) for all models $k = 1, \dots, 5$. The dynamics converges to the homogeneous attractor $S_k = \sum_i O_i / N = 0.3$ so that no categorization takes place: all models fit equally well all data.




Figure 4 – Time evolution of the five modeling fields for fixed fuzziness $\sigma_k = 0.15, \forall k$. The dynamics converges to the local maximum $S_k = 0.3, \forall k$.

In figure 5 we show the results of the same experiment except that the fuzziness is slightly reduced, $\sigma_k = 0.13, \forall k$. This time a symmetry breaking in the space of models takes place, resulting in the emergence of two distinct categories described by the fields $S_k = 3.56$ for $k = 1, 4, 5$ and $S_k = 2.17$ for $k = 2, 3$. As usual, the symmetry-breaking is triggered by inhomogeneities in the initials conditions.




Figure 5 – Time evolution of the five modeling fields for fixed fuzziness $\sigma_k = 0.13, \forall k$. The dynamics converges to the local maximum described in the text.

One might think that decreasing further the fuzziness σ_k will lead to new symmetry-breakings and ultimately to the perfect categorization of all objects. Unfortunately, this is not so: when σ_k is reduced further the partial similarities between concept 1 and all N objects, $l(i|1), \forall i$, become vanishingly small (the argument of the exponential in equation (1) tends to $-\infty$) and hence $f(1|i) \approx 0, \forall i$ so that the modeling field S_1 is never updated. As a result, the system behaves as possessing effectively $M - 1$ adaptive modeling fields. To avoid this

type of difficulty one should always start with large fuzziness to guarantee that at the outset any one model has a nonzero similarity with all objects. Since this choice leads inevitably to the behavior illustrated in figure 4, the solution is to decrease the fuzziness on the flight, i.e., during the time evolution of the modeling fields according to the following prescription

$$\sigma_k^2(t) = \sigma_{k1}^2 \exp(-\alpha t) + \sigma_{k0}^2 \quad (5)$$

with $\alpha = 5 \times 10^{-4}$, $\sigma_{k1} = 1 \forall k$ and $\sigma_{k0} = 0.03 \forall k$. We note that equation (5) differs from the standard MFT formulation [6], but the central idea of updating the fuzziness during the evolution of the modeling fields is the same and constitutes the essence of fuzzy dynamic logic. Application of the standard scheme requires the addition of a mechanism for the elimination of equivalent models and generation of new ones, an issue that we will discuss elsewhere.

In figure 6 we present the results of applying the dynamic fuzziness scheme to the problem of categorizing the five objects discussed above. It is interesting to note that the onset of categorization appears to be associated to the binary splitting of more general concepts, as in the case of the discrimination trees. The reader can easily identify which of the modeling fields converged to a given object by looking at the value of the field at $t=0$ (the list is provided in the paragraph above figure 4).




Figure 6 – Time evolution of the five modeling fields using the dynamic fuzziness scheme. The dynamics converges to the global maximum.

To study how the dynamic fuzziness scheme can adapt the modeling fields when new objects are added to the organism's we repeat the previous experiment using six concept-models (the initial value of this modeling field is $S_6 = 0.82$) but with the same five objects used in the previous experiments.

The results are illustrated in figure 7 and indicate that modeling fields S_1 and S_4 associate to the same object, namely, object 5 for which $O_5 = 0.5$. This is actually a general behavior pattern – categorization is not spoiled by using more concept-models than the number of objects that make up the organism's world. We then repeat the same experiment, but add a new object $O_6 = 0$ at time $t/50 = 200$.




Figure 7 – Time evolution of the six modeling fields using the dynamic fuzziness scheme. The task is to categorize the same five objects of the previous experiments.

The result depicted in figure 8 shows a complete rearrangement of all modeling fields leading ultimately to the perfect categorization of all objects. This success, however, is due to the introduction of the new object at a relatively early stage of the dynamics. If it were introduced at a later stage, say $t/50 = 300$, then only modeling field S_2 would respond by moving towards, and finally fixing at, the mean value $(O_1 + O_6)/2 = 0.5$.




Figure 8 – Time evolution of the six modeling fields using the dynamic fuzziness scheme. The task is to categorize the same five objects of the previous experiments plus a sixth object that entered the world at $t/50 = 200$.

Finally, for the sake of completeness we illustrate in figure 9 the case in which there are more objects (5) than model-concepts (3). This is, perhaps, the situation where the task of categorization is better exemplified since model 1 clumps together objects 4 and 5, model 2 clumps together objects 1 and 2, while model 3 associates to the remaining object 3.




Figure 9 – Time evolution of three modeling fields using the dynamic fuzziness scheme. The task is to categorize the same five objects of the previous experiments.

4. CONCLUSIONS

We have demonstrated the potentiality of the MFT framework, or more precisely its simple variant, as a mechanism for the spontaneous formation of meanings. In contrast to the discrimination tree approach, MFT offers here a genuinely autonomous and efficient mechanism of categorization of objects or situations. An interesting feature of this mechanism, displayed in figures 5 to 9, and which makes it somewhat similar to a discrimination tree, is that the dynamics first merges all modeling fields into a single global model-concept and then proceeds to the refinements through sequential binary divisions. However, further improvements of the present implementation seem to be necessary. Important modifications will include (1) autonomous detection of the number of different objects, (2) efficient handling of new objects added to the organism's world, (3) creating a concept of "object" by differentiating objects from meaningless background, (4) addressing high dimensionality characteristic of visual images and other raw sensory data, and (5) creating abstract concepts, corresponding to situations, rather than simple individual objects. It is possible that the last four modifications could only be addressed by combining meaning creation with communication. Work in these directions is on the way.

ACKNOWLEDGMENTS

Effort sponsored by the Air Force Office of Scientific Research, Air Force Material Command, USAF, under grant number FA8655-04-1-3045. The U. S Government is authorized to reproduce and distribute reprints for Governmental purpose notwithstanding any copyright notation thereon. We are thankful to Ross Deming for his help in this project.

REFERENCES

- [1] C. L. Nehaniv, "The making of meaning in societies: Semiotic and Information-Theoretic background to the evolution of communication," *Proceedings of the AISB Symposium: Starting from Society – the application of social analogies to computational systems*, 73-84, 2000.
- [2] M. Minsky, Plenary address at Artificial Life V, the 5th international workshop on the synthesis and simulation of living systems, 1996.
- [3] L. Steels, "Perceptually grounded meaning creation," *Proceedings of the Second International Conference on Multi-Agent Systems, ICMAS-96*, 338-344, 1996.
- [4] A. D. M. Smith, "Intelligent Meaning Creation in a Clumpy World Helps Communication," *Artificial Life* **9**, 557-574, 2003.
- [5] L. Wittgenstein, *Philosophical Investigations* New York: Basil Blackwell, 1968.
- [6] L. I. Perlovsky, *Neural Networks and Intellect: Using Model-Based Concepts*, Oxford: Oxford University Press, 2001.